Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

In This Issue

To cite this article:
(1987) In This Issue. Operations Research 35(4):486-488. https://doi.org/10.1287/opre.35.4.486

IN THIS ISSUE

Monitoring Advances in Nonlinear Programming Software

The advent of microcomputers, and of spreadsheets and other novel software, as well as advances in mainframe computing, offer many practical opportunities for the operations research community. In parallel with these technological innovations, OR workers in several fields have improved their understanding of computational issues by conducting comparative studies of software and developing guidelines for evaluating and using existing algorithms. In "The Status of Nonlinear Programming Software: An Update," Allan D. Waren, Ming S. Hung, and Leon S. Lasdon report on progress in nonlinear programming software since the 1979 survey featured in this Journal ("The Status of Nonlinear Programming Software" by A. D. Waren and L. S. Lasdon, Vol. 27, pp. 431–456).

Improving the IRS Telephone Taxpayers Information System

In 1986 the Internal Revenue Service received 37.8 million telephone requests for taxpayer assistance, most of which occurred in January (when employers and financial institutions mail income information), and just prior to the April 15th filing date. Even in slower periods, the system is very heavily used: the typical fraction of callers receiving a busy signal may vary from a low of 40% during the nonfiling season to a high of 80% during the filing season. In "Modeling the IRS Telephone Taxpayer Information System," Carl M. Harris, Karla L. Hoffman, and Patsy B. Saunders describe a study, undertaken for the IRS, of this congested environment. The queueing flow balance formula they present is now used by the agency to estimate actual demand at each of its information centers and to make projections that justify budget requests to Congress, as well as to measure network quality of service. The methodology could also be used by the airline industry, which has recently chosen to allow some loss of service during heavy traffic periods in order to reduce the large costs of circuiting and staffing their telephone reservation systems.

An Alternative to "An Eye for an Eye"

Since biblical times, man has debated the question of how to deter aggression. In a noted article, Robert Axelrod and William Hamilton ("The Evolution of Cooperation," *Science*, 1981, Vol. 211, pp. 1390–1396) described a series of computer tournaments whose results supported a "tit-for-tat" strategy in which two opponents simply imitated each other's behavior of the previous round. Axelrod and Hamilton further argued that such strategies of reciprocity are pervasive in our evolutionary history. Now, Steven J. Brams and D. Marc Kilgour, in "Optimal Threats," examine tit-for-tat in a modified situation in which players must decide not only whether to cooperate, but how much to cooperate. They conclude that strict reciprocity may in fact *not* be optimal: more-than-proportionate threats of retaliation against low levels of preemption and less-than-proportionate threats against high levels may be a more effective deterrent in certain conflicts, such as superpower confrontations.

Shoot Now or Later?

In a ballistic missile attack, the defender may wish to use local defenses (short range interceptors) to protect specific targets from approaching missiles (reentry vehicles, or RVs). In responding to such an attack, the defense may have two or more opportunities to intercept attacking missiles. Between these opportunities, it might also be able to see whether it succeeded in destroying the attackers (shoot-look-shoot or even shoot-look-shoot-look-shoot capability). Given these conditions, the defense will seek the best strategy for using the large number of (imperfect) interceptors in its arsenal: for example, should it employ two interceptors per missile for some of the RVs at first, and hope to face a much smaller number in the second round, or should it hold more interceptors in reserve? In "Optimal Terminal Defense Tactics Against Simultaneously Arriving RVs When Several Sequential Engagements Are Possible," Richard M. Soland solves this problem using stochastic dynamic programming. Along the way he shows that, at each opportunity, the interceptors used should be spread as uniformly as possible among the RVs still alive.

Designing a Logistics Network

Why is it that many-to-many transportation networks (by road, air, rail and water) have such varied struc-

tures? Can simple arguments explain why the overnight services of Federal Express and United Parcel Service have each employed a single major hub (in Memphis, Tennessee and Louisville, Kentucky, respectively), but UPS ground service and other trucking networks have many transhipment points? Which key factors influence the network geometry and routing pattern that work best in any given instance? In "The Break-Bulk Role of Terminals in Many-to-Many Logistics Networks," Carlos F. Daganzo answers some of these questions. Examining an idealized model, he develops simple formulas for the (aggregate) network descriptors of a near-optimal system, and for the resulting average origin-to-destination cost. These formulas reveal the ways in which interest rates, vehicle capacity, demand rates and other factors influence network structure. They also are a starting point for more detailed strategic planning analyses.

Scheduling Work Movement in a Flow Shop

In 1985, U.S. manufacturing automation markets were valued at $25.6 billion. This figure is expected to grow by 15% annually to top $50 billion by 1990, and should reach $100 billion in 1995 (*Electronic Business*, November 15, 1985). Increases in factory automation and complexity have heightened the importance of planning for the efficient movement of work within manufacturing facilities. Although researchers in the last few decades have produced an impressive body of literature in deterministic scheduling theory, this community has not developed a comparable amount of knowledge on material handling systems. In "Interstage Transportation Planning in the Deterministic Flow-Shop Environment," Michael A. Langston presents a mathematical model for coordinating flow-shop movement by assessing the delay incurred in both processing tasks and transporting jobs. His analysis yields tight performance guarantees for several fast transportation strategies by combining the notion of compound algorithm design and virtual machine scheduling with a careful sequencing of transportation actions.

Locating Supermarkets in a Competitive Environment

A&P and Safeway have 1,200 and 1,479 stores, respectively, throughout the United States. For many retail businesses, such as supermarkets, store location is crucial to competitiveness: customers' choices of where to shop are based to a large extent on proximity to their residences or workplaces. However, simply locating a business near a large population area may not be enough to ensure success—a competitor might establish its store in an even more convenient location. Thus the successful company must locate its business not only to attract customers, but also to prevent the competition from stealing them. Cost is another important consideration: establishing a retail business is expensive, and thus owners want to minimize the number of locations, subject to preempting the competition. In "Competitive Location on a Network," Gregory Dobson and Uday S. Karmarkar develop a model for such a situation and show how to determine store locations that will be both profitable for the owner and impenetrable by firms that might compete on the basis of location.

Settling into a Steady State

Telephone traffic varies considerably during the day: on an average business day, AT&T customers make 75 million calls, 25% during the peak hours of 10 to 11 A.M. and 3 to 4 P.M. For similar situations, which occur regularly in the contexts of computer operations, communications and manufacturing, analysts often use vacation models to describe the interactions of subsystems with the rest of the system. (See, for example, "A Queue with Starter and a Queue with Vacations: Delay Analysis by Decomposition" by Hanoch Levy and Leonard Kleinrock, in the May–June 1986 issue, and "$D/G/1$ Queues with Vacations" by L. D. Servi, in the July–August 1986 issue, of this Journal.) The dynamic behavior of such systems is of operational interest. Since arrival rates often change with time, steady-state analysis is approximate. One useful measure that helps to justify such approximations is the relaxation time—i.e., the time for the systems to reach a steady state. In "Dynamics of the $M/G/1$ Vacation Model," J. Keilson and L. D. Servi show how to obtain the relaxation time as well as related system performance measures for several situations that arise in practice.

Assessing Queue Lengths as a Job Population Increases

Over the past three decades, queueing networks have become a major tool for modeling and studying large, complex stochastic systems of computers and automated manufacturing. These models (closed queueing networks) often fix the total number of jobs (the "job population") as a constant that typically reflects some aspects of system capacity—for instance, the level of programming in a mainframe computer system, or

the number of pallets in a flexible manufacturing system. In "Stochastic Monotonicity of the Queue Lengths in Closed Queueing Networks," J. George Shanthikumar and David D. Yao study both the dynamic and steady-state behavior of the queue lengths in a processing network when job population increases. This behavior is important, not only because the queue lengths reflect the interdependence of the processors (nodes in the network), but also because the rate of each processor often depends upon the length of the queue it must serve. The authors' results indicate that the rather complex stochastic monotonic behavior of the queue lengths arises essentially because of the deterministic monotonicity of the processing rate at each processor.

Choosing the Clock Rate of Computer Operating Systems

"Clocked operating" real-time computer systems (for example, telephone switches) use an internal clock that ticks every T milliseconds (say, every 10 or every 25 msecs) to schedule software programs. The purpose of this feature is to guarantee timely response to external events (for example, the near-instantaneous dial tone that a caller hears when picking up a telephone receiver). In "Simple Inequalities for the $D/G/1$ Queue," Teunis J. Ott extends the results of earlier investigations in which he constructed a model of such systems using two input streams and translated it into a simple queueing model. Now, he presents a number of inequalities for this queue that can be used to approximately predict the consequences of choosing a particular clock rate, ranging from very simplistic methods (simple enough for use with a pocket calculator) to moderately detailed ones (those that require a simple computer program).

How Long Will the Strike Last?

From 1947 to 1983, employee work stoppages averaged over 284 per year (*Handbook of Labor Statistics, 1985*, U.S. Department of Labor). In assessing the consequences of these strikes and of similar situations that arise in other contexts (for example, length of job durations, reliability of a physical system), practitioners are interested in the distribution of the residual lifetimes of the items under study. Percentiles of the residual lifetimes are a good way to describe these quantitites because they are not very sensitive to extreme values and can be applied before all of the lifetimes are known. "Estimation of Percentage Residual Life" by Miklós Csörgő and Sándor Csörgő presents a nonparametric large sample estimation theory of the percentage residual lifetime. This estimate results in large sample confidence intervals and confidence bands.

Estimating the Flow Capacity of Random Networks

Since the seminal research of L. R. Ford, D. K. Fulkerson and other investigators in the mid-1950s, maximum flow has become a central theoretical and application concept for characterizing a flow network with deterministic arc capabilities. If arc capabilities are random, then the probability distribution function of maximum flow provides the corresponding useful performance measure. Since exact computation of this distribution function is difficult, even for networks of moderate size, the Monte Carlo method can provide an informative estimate of this measure. George S. Fishman, in "The Distribution of Maximum Flow with Applications to Multi-State Reliability Systems," describes and demonstrates such a method. His approach exploits prior information about the network to derive a function estimator with specified accuracy at considerably less cost than would be required to achieve the same accuracy with crude Monte Carlo sampling. The method is designed for arc capacities that diminish by random amounts at random points in time.

Technical Note

In his article, "On The Convexity of Policy Regions in Partially Observed Systems," William S. Lovejoy shows how a simple observation can be used to unify many known convex region results concerning partially observable systems.

A Remembrance of George E. Kimball

George E. Kimball, along with Bernard O. Koopman and Philip M. Morse, was one of the founding fathers of operations research in the United States. ORSA honors his memory by giving his name to its highest award for service to the profession. In "George E. Kimball, A Pioneering Giant in Operations Research," John F. Magee, a former president of ORSA and TIMS and current Chairman and Chief Executive of Arthur D. Little, Inc., recounts some of his experiences in working with Dr. Kimball in the 1950s.